

1. A method, comprising:

providing a plurality of InfiniBand switches, wherein the plurality of InfiniBand

5 switches are coupled to form a Clos network;

providing a plurality of sources coupled to the Clos network;

providing a plurality of destinations coupled to the Clos network;

calculating a plurality of routing trees for the plurality of InfiniBand switches;

10 calculating a plurality of DLIDs and a set of forwarding instructions for each of the plurality of InfiniBand switches, wherein each of the plurality of DLIDs corresponds to one of the plurality of routing trees and one of the plurality of destinations; and

populating a forwarding table of each of the plurality of InfiniBand switches in the Clos network with the plurality of DLIDs and the set of forwarding instructions.

15 2. The method of claim 1, wherein each of the plurality of destinations is identified by a BaseLID.

3. The method of claim 1, wherein calculating the plurality of routing trees comprises for each spine node in the Clos network, calculating a shortest path from the  
20 spine node to each of the plurality of sources and each of the plurality of destinations.

4. The method of claim 1, wherein each of the plurality of routing trees comprises at least a portion of the plurality of InfiniBand switches and corresponding plurality of links that form a shortest path from one of the plurality of sources or one of the plurality of  
25 destinations to a spine node of the Clos network.

5. The method of claim 1, further comprising:

creating a packet at one of the plurality of sources, wherein the packet is addressed to one of the plurality of destinations;

30 executing a rearrangement algorithm for the Clos network;

assigning one of the plurality of DLIDs to the packet; and

the packet following a path through at least a portion of the plurality of InfiniBand switches from the one of the plurality of sources to the one of the plurality of destinations,

wherein each of the portion of the plurality of InfiniBand switches forward the packet according to the one of the plurality of DLIDs assigned to the packet.

5       6. The method of claim 5, wherein the Clos network operates as a strictly non-interfering network.

7. The method of claim 5, wherein the packet following the path comprises looking up the one of the plurality of DLIDs assigned to the packet in the forwarding table at each of the portion of the plurality of InfiniBand switches along the path from the one of  
10       the plurality of sources to the one of the plurality of destinations.

8. The method of claim 5, wherein calculating the plurality of routing trees comprises calculating the plurality of routing trees sufficient to execute the rearrangement algorithm.

15       9. The method of claim 5; wherein the packet following the path comprises each of the portion of the plurality of InfiniBand switches forwarding the packet in accordance with the one of the plurality of DLIDs assigned to the packet as found in the forwarding table at each the portion of the plurality of InfiniBand switches.

20       10. A method, comprising:  
providing a plurality of InfiniBand switches coupling a plurality end nodes to form a network;

calculating a plurality of routing trees for the plurality of InfiniBand switches;  
25       calculating a plurality of DLIDs and a set of forwarding instructions for each of the plurality of InfiniBand switches, wherein each of the plurality of DLIDs corresponds to one of the plurality of routing trees and one of the plurality of end nodes; and  
populating a forwarding table of each of the plurality of InfiniBand switches in the network with the plurality of DLIDs and the set of forwarding instructions.

30       11. The method of claim 10, wherein the network is a Clos network.

12. The method of claim 10, wherein each of the plurality of end nodes comprises a destination, and wherein the destination is identified by a BaseLID.

13. The method of claim 10, wherein calculating the plurality of routing trees  
5 comprises for each spine node in the network, calculating a shortest path from the spine node to each of the plurality of end nodes.

14. The method of claim 10, wherein each of the plurality of routing trees  
comprises at least a portion of the plurality of InfiniBand switches and corresponding  
10 plurality of links that form a shortest path from one of the plurality of end nodes to a spine node of the network.

15. The method of claim 14, wherein the shortest path is loop-less.

16. The method of claim 10, further comprising while the network is in operation,  
transforming the network into a Clos network, and wherein after transforming to the Clos  
network, repeating the method of claim 11 on the Clos network.

17. A method, comprising:  
20 providing a plurality of InfiniBand switches coupling a plurality sources and a plurality of destinations to form a Clos network;  
creating a packet at one of the plurality of sources, wherein the packet is addressed to one of a plurality of destinations;  
executing a rearrangement algorithm for the Clos network;  
25 assigning one of a plurality of DLIDs to the packet; and  
the packet following a path through at least a portion of a plurality of InfiniBand switches from the one of the plurality of sources to the one of the plurality of the destinations, wherein each of the portion of the plurality of InfiniBand switches forward the packet according to the one of the plurality of DLIDs assigned to the packet.

30

18. The method of claim 17, wherein the Clos network operates as a strictly non-interfering network.

19. The method of claim 17, wherein the packet following the path comprises looking up the one of the plurality of DLIDs assigned to the packet in a forwarding table at each of the portion of the plurality of InfiniBand switches along the path from the one of the plurality of source to the one of the plurality of destinations.

5

20. The method of claim 17, wherein the packet following the path comprises each of the portion of the plurality of InfiniBand switches forwarding the packet in accordance with the one of the plurality of DLIDs assigned to the packet as found in a forwarding table at each the portion of the plurality of InfiniBand switches.

10